




## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : <b>G06F 12/14</b>		A1	(11) International Publication Number: <b>WO 97/36235</b>
			(43) International Publication Date: <b>2 October 1997 (02.10.97)</b>
<p>(21) International Application Number: <b>PCT/IB97/00230</b></p> <p>(22) International Filing Date: <b>10 March 1997 (10.03.97)</b></p> <p>(30) Priority Data: 08/621,759 22 March 1996 (22.03.96) US</p> <p>(71) Applicant: <b>PHILIPS ELECTRONICS N.V. [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).</b></p> <p>(71) Applicant (for SE only): <b>PHILIPS NORDEN AB [SE/SE]; Kottbygatan 7, Kista, S-164 85 Stockholm (SE).</b></p> <p>(72) Inventors: <b>WENDORF, James, W.; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). RATH, Kamlesh; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). VERMA, Dinesh; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).</b></p> <p>(74) Agent: <b>GROENENDAAL, Antonius, W., M.; Internationaal Octrooibureau B.V., P.O. Box 220, NL-5600 AE Eindhoven (NL).</b></p>			<p>(81) Designated States: <b>JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).</b></p> <p><b>Published</b>  <i>With international search report.  Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>

(54) Title: PROTECTION DOMAINS IN A SINGLE ADDRESS SPACE

(57) Abstract

Protection among threads executing in the same address space of a computer system is provided without using virtual memory techniques. This is achieved by grouping the threads into protection domains, each of the threads in a protection domain having the same rights to access memory as the other threads in that protection domain, so that each thread in a protection domain can access all the information available to the others. At least one protection domain, referred to herein as the "system" domain, which typically is the protection domain of the operating system and has unrestricted access to the entire memory, is predefined prior to execution of any threads. Prior to execution, the single address space is divided into non-overlapping pages. Each page has at least one access permission set for it. Only threads that belong to a protection domain having permission to access a page may do so. During operation, when a request to access memory is issued by an executing thread, it is determined whether or not the protection domain of the executing thread has permission to perform the requested type of access. If the protection domain of the executing thread is permitted to perform the type of access requested, access is granted and the executing thread's execution proceeds normally. However, if the protection domain of the executing thread does not have permission to perform the requested type of access, a protection fault is generated.



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	N	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	R	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

Protection domains in a single address space.

### Technical Field

This invention relates to the field of computer operations, and in particular, to reducing the likelihood of interference by one process with the proper operation of another process.

5

### Background of the Invention

Computer operating systems such as UNIX provide a sophisticated degree of protection among processes concurrently executing on the system. A process is defined as a program in execution which utilizes system resources, such as memory and computing

10 time. At each step of execution the process generates the address, or addresses, in memory which is needed to successfully execute the step, e.g., a) the address where an instruction to be executed is stored, b) the address where required data is to be found, or c) the address where output data is to be stored. The range of addresses that a computer can generate defines its address space.

15 Under control of a conventional operating system, e.g., UNIX, the addresses generated by a process do not directly address a location in the computer's physical memory. Instead, the addresses generated by the process are translated to a physical location using well known virtual memory techniques. A description of virtual memory can be found in the book "Operating Systems Concepts", by Peterson and Silbersatz. Each process

20 executes its steps and generates addresses as if it is executing on its own private computer which has a very large memory. The size of the memory on this virtual private computer is the size of the process' address space. The addresses generated by the process in this virtual memory are translated to physical memory locations, also known as real memory addresses, using a page table mechanism, typically implemented in a memory management unit (MMU).  
25 Thus, each process is provided a different address space, which is mapped onto the real memory. The MMU translations are briefly described below.

The physical memory is divided into fixed units called pages. The page size is determined by the hardware architecture of the processor. The address space of each process is also divided into pages, each having the same size as a page of the physical

memory. For each process, a page table is maintained in memory by the operating system which contains a mapping of the page number in the address space to the page number in the physical memory. The page table also contains some attributes bits for each page. The attribute bits indicate a) whether the page with which they are associated is valid or invalid, 5 i.e., whether or not the associated page has a corresponding page in physical memory, and b) if the process has permission to read or to modify the contents of the page.

When an address is generated by a process, the MMU first looks up the page table in the memory and determines the corresponding physical page in the memory. The permissions for the page are also checked during the look up process to verify that the 10 process is authorized to access the memory in the manner, e.g., read or write, that it is attempting. The physical page number is attached to the offset of the address within the page and a new physical address is generated. The physical address is then used to look up the contents of the memory. Thus, each memory access translates into two sequential accesses, the first one for generating the physical address and checking the validity, with the second 15 one being for actually accessing the physical memory.

The two accesses are usually expedited by means of caching techniques, in which frequently used areas of the page tables (and memory) are stored in a smaller higher speed memory called a cache so as to reduce the probability of actually accessing the physical memory.

20 One of the advantages of going through a page table is that processes can execute relatively independent of each other. A process has exclusive access to its address space and no other process can modify the contents of the address space, unless the process explicitly provides access to parts of its address space to other processes.

In a system with multiple concurrent processes, switching between 25 processes is a complex, and so "expensive" in processing terms, operation. A process may be further broken down into several threads, each thread being a sequence of executing instructions. All threads within a process execute in the same address space. As such, switching from the currently active thread to another thread is a relatively simple and inexpensive operation. However, all threads have identical privileges to all of the address 30 space of the process to which they belong, and so there is no protection among threads executing in the same address space.

In the area of embedded systems, such as 1) digital television receivers, 2) television set top units, and 3) network switch controllers, there is only one process in the entire system. This process may consist of multiple concurrent threads. However, typically.

there is no support for virtual memory because such systems have only so-called "primary memory", e.g., RAM, ROM or combination thereof, and no so-called "secondary storage", e.g., a disk. Thus, all the threads have access to the entire memory, including data structures maintained by the operating system.

- 5                   This global access is a serious problem in the development and debugging of embedded system software. For example, an errant thread, e.g., one still under development, can alter a data structure or code used by an already debugged and "trusted" thread. The trusted code will eventually be affected by the corruption of its data structure or code, and either generate an error or behave unexpectedly. It is very difficult to determine
- 10                  the cause of the problem in such an environment, since the point in the code at which the software crashes or generates a warning is unrelated to the point in the code that caused the problem.

An identical problem exists in digital television receivers which are used for the downloading of multiple applications from a server. The downloaded application may cause the receiver to crash, and it is difficult to isolate the cause of the fault, which could lie in the downloaded application software, the operating system software, or any other supporting library. When the software in the embedded system originates from multiple program development organizations that are cooperating, it is difficult to isolate the cause of the problem. As a result, each organization typically blames the other for the failure.

- 20                  In embedded systems where an MMU is present, groups of threads can be logically clustered into processes, which are each provided their own address space using MMU page mapping. In this manner, different threads can be provided isolation from one another. In some systems, the MMU is preloaded so that the physical page numbers correspond to the virtual page numbers and the physically generated address is identical to the virtual address.
- 25                  Unfortunately, the cost of having an MMU is usually unacceptable in an embedded system. In those cases, there is no protection provided among the threads.

#### Summary of the Invention

We have realized that, from both the debugging and development perspective, providing protection among threads executing in the same address space is an attractive idea. Therefore, in accordance with the principles of the invention, protection among threads executing in the same address space of a computer system is provided without using virtual memory techniques that require each thread actually, or logically, to be isolated in its own, separate address space. This is achieved by grouping the threads into protection

domains, each of the threads in a protection domain having the same rights to access memory as the other threads in that protection domain, so that each thread in a protection domain can access all the information available to the others. At least one protection domain, referred to herein as the "system" domain, which typically is the protection domain of the operating system and has unrestricted access to the entire memory, is predefined prior to execution of any threads. The protection domains, including the system domain, may spawn additional protection domains. However, the total number of protection domains is typically limited to a predetermined maximum number. Also prior to execution, the single address space, which is typically physical memory only and is common to all of the threads, is divided into pages, which typically do not overlap. Each page, either prior to or during execution, has at least one access permission set for it, e.g., on a protection domain-by-protection domain basis or on an exception basis. Only threads that belong to a protection domain having permission to access a page may do so. Permission for read access and write access to each page may be separately specified.

15 During operation, when a request to access memory is issued by an executing thread, it is determined whether or not the protection domain of the executing thread has permission to perform the requested type of access. If the protection domain of the executing thread is permitted to perform the type of access requested, access is granted and the executing thread's execution proceeds normally. However, if the protection domain 20 of the executing thread does not have permission to perform the requested type of access, a protection fault is generated.

It is preferable to identify the protection domain of the currently executing thread prior to the memory access. Such identification may be performed each time memory is accessed or each time the thread is changed.

25

#### Brief Description of the Drawing

In the drawing:

FIG. 1 shows a portion of the hardware architecture of a processor for use in practising the invention;

30 FIG. 2 shows a flow chart of an exemplary process for setting the value of the protection domain register when there is a switch from one thread to another in accordance with an aspect of the invention;

FIG. 3 shows an exemplary process for determining whether a particular access to memory is valid, in accordance with an aspect of the invention;

FIG. 4 shows an exemplary process for the handling of an operating system call by a currently executing thread, in accordance with an aspect of the invention;

FIG. 5 shows an exemplary process by which memory is allocated for a conventional operating system call in a system employing protection domains, in accordance 5 with the principles of the invention; and

FIG. 6 shows an exemplary process by which a new protection domain is created, in accordance with an aspect of the invention.

#### Detailed Description

10 FIG. 1 shows a portion of the hardware architecture of a processor for use in practising the invention. As shown in FIG. 1, the portion of the architecture to the left of dashed line 117 is conventionally found in processors and is well known in the art while the portion of the architecture to the right of dashed line 117 is specific to the implementation of the invention. The architecture of FIG. 1 includes: a) general purpose 15 registers 101, b) instruction decoder 103, c) execution unit 105, d) memory address register 107, e) access list table 109, f) protection domain register (PDR) 111, e) read/write (R/W) 113, and f) domain access list 115.

General purpose registers 101 include general purpose registers 101-1 through 101-N. Each of general purpose registers 101 is used by the processor for storing 20 data, typically on a temporary basis, and executing instructions that require the data. Also, register 101-1 is the program counter (PC), which is used to point to the address of the next instruction to be executed while register 101-2 is the stack pointer (SP), which is used to point to locations at which context information is stored during subroutine execution. Each 25 thread may maintain its own stack which, while the thread is executing, is pointed to by the thread's own stack pointer, which may be implemented using register 101-2.

Instruction decoder 103 receives instructions from memory (not shown) and decodes the instructions in the conventional manner. Instruction decoder 103 is coupled to execution unit 105 in order to supply control and other information to execution unit 105 for executing the received instructions. Execution unit 105 is coupled to general purpose 30 registers 101 for manipulating data therein, including: 1) writing data to general purpose registers 101, 2) receiving information back from general purpose registers 101, and 3) performing other operations on the contents of general purpose registers 101.

Memory address register 107 stores the address of a memory request prior to that address being supplied to the address bus so that the validity of the memory access

can be determined. Each memory address is arranged as a page number and an offset within the page, in the conventional manner. The number of a particular page being accessed is stored in page number 107-1 and the corresponding offset information is stored in offset 107-

2. All of the addresses are within the same address space. Thus, identical memory  
5 addresses used by different threads, even when the different threads are in different protection domains, always correspond to the same physical memory location.

Access list table 109 contains the read/write access permissions for each page. In particular, each of the entries 109-1 through 109-N specifies the read/write permissions for one page of memory. Table 1 shows an exemplary embodiment of access  
10 list table 109 as shown in Table 1. The access list table is arranged so that its row index is the page number and its columns are indexed by a double index including the protection domain as primary index and whether the access is a read or a write as the secondary index. In other embodiments of the invention the read or write control may be unified so that it is identical for each protection domain, and, therefore, only one column need be stored for  
15 each protection domain.

In other embodiments of the invention, the access list table might group several pages into a common memory unit, so that all pages in the same memory unit have identical entries in the access list table.

Table 1

20

Protection Domain	(System)		1		..	M	
	0		1			M	
Page #	R	W	R	W	..	R	W
0	1	1	0	0	..	1	1
1	1	1	1	1	..	1	1
2	1	1	1	1	..	1	1
3	1	1	0	0	..	0	0
:	:	:	:	:		:	:
N	1	1	0	0	..	0	0

30

Returning to FIG. 1, in the embodiment of the invention shown therein, the page number of the memory address to be accessed from page number 107-1 is used to index the pointer into access list table 109 and to determine a particular row of permissions. This row is then copied into domain access list 115 a register which stores one row of access list table 109.

Each time there is a context switch, i.e., another thread begins executing, protection domain register 111 is loaded with a value indicative of the protection domain of the new thread that is beginning to execute. Doing so is handled by the operating system, which has information indicative of the thread that is executing and the protection domain to which it belongs. In particular, the operating system conventionally maintains the identity of the currently active thread in a fixed memory location or register. Associated with each thread, the operating system conventionally maintains a thread-control block, which is a portion of memory containing information pertinent to the thread, and used by the operating system. The protection domain to which the thread belongs can be stored in the thread-control block, from where it can be readily accessed by the operating system. An alternate scheme would be to include the protection domain as part of the thread identifier, for example, the first few bits of a thread-identifier can indicate the correct protection domain.

Read/write flag 113 is loaded with a value indicative of whether a particular memory access is to be a read or a write. For example, a 0 can indicate a read and a 1 can indicate a write.

Once domain access lists 115 is loaded with the particular row from access list table 109 that corresponds to the page that is being accessed, 1) the value stored in PDR 111, which is indicative of the protection domain that is currently executing, and 2) the value of read/write 113, which indicates the type of access requested, are used to index into the value stored in domain access list 115. If the protection domain that is executing is entitled to perform the requested access, the value obtained from domain access list 115 will indicate that the access is valid, e.g., a logic 1. Otherwise, the value retrieved from domain access list will indicate that the access is invalid, e.g., a logic 0. Should an invalid access be detected a protection fault is generated, in accordance with the principles of the invention.

In another embodiment of the invention, a protection domain's access permissions may be determined by first indexing into access list table 109 using the value indicative of the current protection domain, which is stored in PDR 111, to extract the current protection domain's column of permissions for each memory page, and then using the

value of page number 107-1 to determine the permissions for each particular page when it is accessed. Doing so allows a protection domain's column of permissions for each memory page to be extracted only when the protection domain is changed, rather than for every memory access, which results in greater speed. Also, advantageously, the columns that correspond to access list table 109 can be stored at independent locations in memory. Thus, the values indicative of the protection domains, e.g., which are loaded in PDR 111, can be the addresses of the memory locations at which is stored the beginning of each of the columns of access list table 109.

FIG. 2 shows a flow chart of a exemplary process for setting the value of the protection domain register when there is a context switch, i.e., a switch from one thread to another, in accordance with an aspect of the invention. In particular the process is entered in step 201, upon the occurrence of a context switch, e.g., the time period allotted to a particular thread that has been executing is over and, therefore, the operating system stops execution of that particular thread and initiates the execution of another thread. Next, conditional branch point 203 tests to determine if the thread whose execution is being initiated is in the protection domain of the operating system. If the test result in step 203 is YES, control passes to step 205, in which the operating system sets the PDR to indicate system. The process then exits in step 209. If the test result in step 203 is NO, control passes to step 207, in which the PDR is set to a value indicative of the protection domain of the new thread. The process then exits in step 209. According to the embodiment of the invention described above in which access list table 109 is indexed into first by column, the value indicative of the protection domain to which the PDR is set in step 207 is the first address of the column of access list table 109 for the protection domain.

FIG. 3 shows an exemplary process for determining whether a particular access to memory is valid, in accordance with an aspect of the invention. The process is entered in step 301, when a thread attempts to access a memory location. Conditional branch point 303 tests to determine if the currently executing thread is in the protection domain of the operating system. If the test result in step 303 is YES, the access is valid because the protection domain of the operating system encompasses the entire memory and has access to all of memory. Therefore, control passes to step 305 and the access proceeds as requested. The process then exits in step 307.

If the test result in step 303 is NO, control passes to step 309, in which the page number is extracted from the address supplied by the thread for the memory access. Next, in step 311, the domain access list is determined by reading from the access list table

the value pointed to by the page number in which the memory access is requested. Thereafter, conditional branch point 313 tests to determine if the current domain has access to the requested page, e.g., by checking the value of the bit in the domain access list that corresponds to the type of access requested for the domain of the executing thread. If the 5 test result in step 313 is YES, the memory access is valid and permitted. Therefore, control passes to step 305, and the process continues as described above. If the test result in step 313 is NO, the access is not permitted and is invalid. Therefore, control passes to step 315, in which a protection fault is generated. The process then exits in step 307.

FIG. 4 shows an exemplary process, in accordance with an aspect of the 10 invention, for handling the issuance, by a currently executing thread, of an operating system call, i.e., a request by an executing thread for a service provided by the operating system, e.g., writing particular information to a specified display device.

The process is entered in step 401 upon a call by the currently executing 15 thread to the operating system for service. In step 403, the value of the protection domain register is stored into a predetermined location in memory, e.g., on the stack of the currently executing thread, the next available address of which is pointed to by the thread's stack pointer. Thereafter, in step 405, the value of the protection domain register is changed from the value of the executing thread which is issuing the operating system call to the value of 20 system, as described above. Next, in step 407, a traditional real time operating system call is performed. For example, the context is switched to the operating system and the particular information that had been specified by the thread is written to and displayed on a screen by the operating system. As part of completing the call, the context is restored to the previously executing thread by the operating system, in the conventional manner.

Upon completion of the operating system call control passes to step 409, 25 in which the value of the previously stored protection domain register is retrieved, e.g., popped from the stack, and restored in the PDR, in accordance with an aspect of the invention. Thereafter the process exits in step 411.

The system described above is further assisted by an operating system 30 interface that permits applications to create and manipulate protection domains. In general, conventional operating system interfaces provide several operating system calls, which, when invoked by an application, perform specific, predefined tasks for that application. These tasks often include a) creation and/or manipulation of a thread, b) facilitation of inter-thread communication, and c) creation, use and modification of inter-thread synchronization objects, e.g., message queues, semaphores, or the like.

Some objects may be created for the exclusive use of threads in a particular protection domain. Such objects are created in pages for which only the object creating the protection domain has permission to access. Alternatively, other objects may be created so that they can be used by threads in the protection domain of the creating thread as well as threads in other protection domains, e.g., a message queue by which threads in one protection domain send requests to threads in another protection domain. Such objects are created in a memory page that is a primary memory page of the protection domain of the creating thread but for which one or more other protection domains have permission to access that page, e.g., according to the access list table.

10 All objects created by the system must be stored somewhere in memory. In particular, objects must be stored in memory pages that are accessible by the thread, or threads, which require and can manipulate them. One method that can be used to select the pages of memory in which to create an object is an augmented operating system call. In particular, the operating system call is augmented so as to indicate the protection domains 15 which can have access to the object being created. This approach has the drawback of not being backward compatible, in that software developed prior to implementation of the protection domain system can not be reused.

In accordance with an aspect of the invention, new operating system calls are added which implement protection domains in a manner that is transparent to existing 20 software. Advantageously, existing operating system calls need not be changed. To achieve this, the operating system associates a set of memory pages with each protection domain. Any executing thread of a protection domain at least has access to the primary memory pages 25 associated with the protection domain to which the thread belongs. When a conventional operating system call that requires memory is made by a thread, the memory required for the operating system call is allocated from the primary pages that are associated with the protection domain of the thread.

Each protection domain may be able to access other memory pages that are not associated therewith as a primary memory page provided it has permission to do so. The operating system also contains a function that permits a thread to change the set of 30 primary pages for its protection domain. However, a thread can change the primary pages of its protection domain, e.g., via the operating system, only to those pages that the protection domain has both read and write access permission. Only a protection-domain-aware thread is capable of doing so.

Table 2 shows an exemplary primary pages table. The data shown within

the table is strictly for pedagogic purposes to give a feel for one arrangement of the table. Note that, at the discretion of the implementor, the same memory page may be associated as a primary page for more than one protection domain.

5

Table 2

Protection Domain	Primary Pages
0 (System)	0, 1, 2, 5, ..., M
1	8, 9, 10, 15, 16, 22
:	:
N	3, 7, 11, 82, 89, 105

FIG. 5 shows an exemplary process by which memory is allocated for a conventional operating system call when protection domains are employed, in accordance with the principles of the invention. The process is entered, in step 501, when a request for an allocation of memory is received from an executing thread. Conditional branch point 503 tests the value of the PDR to determine if the requesting thread is in the protection domain of the operating system. If the test result in step 503 is YES, control passes to step 505, in which the required memory is allocated from the available free pages, e.g., those pages not assigned as primary to anyone or those pages which are unused but are assigned as primary to the system protection domain. The process is then exited in step 507.

If the test result in step 503 is NO, indicating that the requesting thread is not in the protection domain of the operating system, conditional branch point 509 tests to determine if allocating the extra memory would cause the resulting total amount of memory allocated to the protection domain of the requesting thread to exceed the amount of memory available to that protection domain's primary pages. If the test result in step 509 is YES, indicating that allocating the extra memory would cause the available memory to be exceeded, the memory allocation is denied in step 511. The process is then exited in step 507. If the test result in step 509 is NO, indicating that the memory available to the protection domain of the thread will not be exceeded by allocating the requested memory, then the particular primary pages of memory for the thread's protection domain are determined in step 513, and the

requested memory is allocated therefrom in step 515. The process then exits in step 507.

In other embodiments of the invention, memory may be pre-allocated for segments from the system memory, or pages may be brought into the primary set on a demand-basis from the system memory. In still other embodiments of the invention, all threads

- 5 can only access memory locations that are within memory pages that are primary pages for the protection domain of the thread. For non-protection-domain-aware threads such an embodiment of the invention operates the same as described above. However, operation is different for protection-domain-aware threads. Such protection-domain-aware threads that require access to memory locations that are not primary pages for the thread's protection domain must first
- 10 instruct the operating system to make any such pages primary for the thread when it requires access.

FIG. 6 shows an exemplary process by which a new protection domain is created, in accordance with an aspect of the invention. This process is typically performed only by protection-domain-aware threads. The protection domain creation process is entered in step 15 601 upon the issuance of an operating system call by a thread requesting the creation of a protection domain with certain attributes. For example, the operating system call may specify the number of primary memory pages to be associated with the protection domain to be created, as well as the location and attributes of the first thread to be created therein. In step 603, the current value of the PDR is stored for later use, e.g., on the creating thread's stack, and in step 20 605 the value of the PDR is changed to that of system.

Next, in step 607, the system assigns a new identifier, i.e., a value that can be stored in the PDR, for use in identifying the protection domain being created. In step 609, the system determines the initial set of primary pages to be assigned to the new protection domain. Thereafter, in step 611, the permissions for the initial set of primary pages of the new 25 protection domain are set in the access list table. In step 613, the new domain identifier is loaded into the PDR, thus designating the newly created protection domain as the currently executing protection domain. Subsequently, in step 615, the first thread of the newly created protection domain begins executing. Steps 613 and 615 may be performed by storing the new domain identifier on the stack of the first thread of the newly created protection domain and then 30 performing a context switch from the operating system to that first thread.

When the creating thread begins execution again, the previously stored value of the old PDR is retrieved, e.g., it is popped off its stack, and replaced in the PDR in step 617. Eventually, the process exits in step 619.

Those of ordinary skill in the art will recognize that the foregoing process

may be arranged according to the requirements of the application. For example, the creating thread may continue to execute after creation of the new protection domain and a thread of the new protection domain only later begins execution when certain predefined conditions are met.

In other embodiments of the invention, association of primary pages may be  
5 done, fully or partly, on a per-thread basis.

Given the foregoing, in addition to creating protection domains, those of ordinary skill in the art will be able to develop processes that manipulate and destroy protection domains.

The foregoing merely illustrates the principles of the invention. It will thus  
10 be appreciated that those skilled in the art will be able to devise various arrangements which, although not explicitly described or shown herein, embody the principles of the invention and are thus within its spirit and scope.

CLAIMS:

1. A method for use in controlling memory access by a plurality of threads which execute in a single address space of a computer system, the method comprising the steps of:

grouping the threads into protection domains, each of the threads in a particular

5 protection domain having the same rights to access the memory of the single address space as the other threads in that particular protection domain so that each thread in a particular protection domain is permitted to access all memory addresses that can be accessed by the other threads in the same particular protection domain;

dividing memory into pages, each page being (i) within the single address space

10 which is common to the threads and (ii) having at least one access permission set therefor so that only threads in a protection domain having access permission can validly access the page;

receiving, from a thread in a currently executing protection domain, a request for a particular type of access to a memory location in a particular page; and

15 generating a protection fault if the currently executing protection domain does not have access permission for the particular type of access requested for the particular type of access to the particular page.

2. The invention as defined in claim 1 further including the steps of determining a currently executing protection domain.

3. A computer processor for executing threads in a single address space,

20 comprising:

a memory page address register for storing a page identifier for a memory address to which access is requested by a currently executing one of said threads; and

a memory page permission store for storing access permissions;

25 wherein a protection fault is generated substantially concurrently with a request by the currently executing one of the threads for access to a memory location identified by a particular page identifier placed in the page address register for which the memory page permission store does not indicate that a valid permission is stored for the protection domain of the currently executing one of the threads.

4. The invention as defined in claim 3 wherein the microprocessor further

comprises a protection domain register for storing an indication of the protection domain of the currently executing one of the threads.

5. The invention as defined in claim 3 or 4 wherein the memory page permission store stores independent access permissions for at least one of said pages for reading and writing operations and the microprocessor further comprises a read/write register for storing the type of access requested by the currently executing one of the threads and for supplying an indication thereof to the memory page permission store.

6. The invention as defined in claim 3, 4 or 5 wherein the memory page permission store stores the access permissions on a protection-domain-by-protection-domain basis.

7. The invention as defined in claim 6 wherein the memory page permission store is arranged to store a read access permission and a write access permission for at least one of the protection domains.

8. A computer processor for executing threads in a single address space, said address space being divided into memory pages, the computer processor comprising:

means for storing page-by-page memory page access permissions as a function of protection domains into which threads which may execute on said processor are grouped;

means for indexing into said access permissions to determine if a particular protection domain of a particular executing thread which is requesting access to a particular memory address has access permission for the page in which the particular memory access is located; and

means for generating a protection fault substantially concurrently with said access request by the particular executing one of the threads for which the stored access permissions do not indicate a valid permission for the protection domain of the particular requesting thread.

25 9. The invention as defined in claim 8 wherein the processor further comprises means for accessing the particular memory address when the stored access permissions indicate a valid permission for the protection domain of the particular requesting thread.

10. The invention as defined in claim 8 or 9 wherein the means for indexing comprises a page number portion of a memory access register.

30 11. The invention as defined in claim 8, 9 or 10 wherein the means for indexing comprises a protection domain register.

12. The invention as defined in claim 8, 9, 10 or 11 wherein the means for indexing includes a read/write flag.

13. The invention as defined in claim 8, 9, 10, 11 or 12 wherein the processor

further includes means for storing its operating context when processing an operating system call.

14. The invention as defined in claim 8, 9, 10, 11, 12 or 13 wherein the means for storing page-by-page memory page access permissions stores a read access permission and  
5 a write access permission for at least one memory page.

1/6




FIG. 1

2/6




FIG. 2

3/6




FIG. 3

4/6




FIG. 4

5/6




FIG. 5

6/6




FIG. 6

1  
INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB 97/00230

## A. CLASSIFICATION OF SUBJECT MATTER

IPC6: G06F 12/14

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC6: G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE, DK, FI, NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WPI

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GB 2260004 A (APPLE COMPUTER INC), 31 March 1993 (31.03.93), page 3, line 22 - page 4, line 11; page 26, line 1 - page 27, line 23, figures 1,6  -- -----	1-5,7-12,14

 Further documents are listed in the continuation of Box C. See patent family annex.

- \* Special categories of cited documents:
- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search

18 August 1997

Date of mailing of the international search report

21-08-1997

Name and mailing address of the ISA/  
Swedish Patent Office  
Box 5055, S-102 42 STOCKHOLM  
Facsimile No. +46 8 666 02 86Authorized officer  
Jan Silfverling  
Telephone No. +46 8 782 25 00

**INTERNATIONAL SEARCH REPORT**  
Information on patent family members

06/08/97

International application No.  
PCT/IB 97/00230

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
GB 2260004 A	31/03/93	JP 5233458 A US 5581722 A	10/09/93 03/12/96